

(11) Publication number:

0 273 083
A1

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 86430058.7

(51) Int. Cl.4: G06F 5/06

(22) Date of filing: 30.12.86

(43) Date of publication of application:
06.07.88 Bulletin 88/27

(64) Designated Contracting States:
DE FR GB IT

(71) Applicant: International Business Machines Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)

(72) Inventor: Giroir, Didier Francis Villa "Le Mazet"
Groupe les Genêts Domaine des Semboules
F-06600 Antibes(FR)
Inventor: Mullery, Alvin Paul
29, Chemin du Lautin
F-06800 Cagnes-sur-Mer(FR)
Inventor: Pauporte, André
Lotissement du Colombier
F-06580 La Colle sur Loup(FR)

(74) Representative: Bonin, Jean-Jacques
Compagnie IBM France Département de
Propriété Intellectuelle
F-06610 La Gaude(FR)

(54) Non-locking queueing mechanism.

(57) Non-locking queueing mechanism for transferring information from a sending unit (100) to a receiving unit (110) through a queue (120) in which there is no interference between the independent units (sender 100 and receiver 110), while enqueueing or dequeuing and thus avoiding any form of interlock or serialism. The mechanism includes a first pointer (D) identifying the element area in said queueing means (120) where the last dequeued information element, if any, was located, a second pointer register (640) for logging a second pointer (E) identifying the element area in said queueing means where the last enqueued information element, if any, was located, a first control block (610) activated by said sending unit (100) to enqueue said information element into said queueing means (120) and for updating said second pointer, and a second control block (620) activated by said receiving means to dequeue said information element from said queueing means (120) and for updating said first pointer.

EP 0 273 083 A1




FIG. 6

NON-LOCKING QUEUEING MECHANISM

Technical Field

This invention relates to digital communication systems and, in particular, to a queuing mechanism enabling a receiver device to be decoupled from a sending device.

Background of the Invention

Most communications between independent hardware components of a computer installation (from very simple units to sophisticated processing systems) takes place by a complex protocol insuring that the sending and receiving hardware components are synchronized. A channel is an example of such a communication means. This need for synchronization means that the receiving component must interrupt the work that it is performing, often at great cost to save and restore the status of the work.

A better alternate is to have some form of queuing mechanism, enabling the receiver to be decoupled from the sender. With a queue, a sender can transmit the data to the receiver's queue (enqueue) while the receiving hardware continues to do its normal function. At its own convenience, the receiving hardware can actually receive (dequeue) the transmitted data. Since the sending and receiving hardware components are independent entities capable of acting in parallel, and since the queue is, in effect, a shared resource, a lock normally is necessary to prevent the sender and receiver from interfering with each other during the enqueue and dequeue. This reduces the level of parallelism and the performance and also increases the complexity and the amount of information to be exchanged.

Generally, the lock between units, say A and B is implemented:

-explicitly via Test And Set or Compare And Swap instructions (or any other kind of hardware or software implementation of semaphores), if both A and B units are processors, to monopolize the queue during an enqueue or a dequeue manipulation, or

-implicitly, if both units use a single "Queue Manipulation Hardware" whose purpose is execute the Enqueue and Dequeue commands received from the units. Since this mechanism can execute only one command at a time, a serialism is done. The Lock is then implicitly implemented via this form of serialism.

U.S. Patent 4,507,760 shows a FIFO memory configuration in which the lock is explicitly implemented. When both writing information into the FIFO queue and reading information from the FIFO queue are requested, a control circuit 26 gives priority to the write operation which eventually locks the communication between the queue and the receiving device.

10

Summary of the Invention

In accordance with the illustrative embodiments of the present invention, a special queuing and dequeuing mechanism is disclosed in which there is no interference between the independent hardware entities (sender and receiver), while enqueueing or dequeuing and even in the generation and propagation of status signals when needed. Therefore, for the mechanism presently disclosed, no interlocking (implicit or explicit) is required. The invention involves a storage device which can be random access memory for storing N work elements (WE) where N is an integer. The control of the memory uses two pointer registers E and D. D pointer points to the last WE dequeued from the queue. E pointer points to the last WE enqueued in this queue.

When information is to be entered in storage for enqueueing, a test is performed to check whether the queue is full or not. In the latter case, the WE is stored in Storage at the address INCR(E) and E pointer is incremented where INCR points to the location logically following that pointed to by E.

When a WE is to be read from the queue, a test is effected to ensure that the queue is not empty. In this case a WE is dequeued from address INCR(D) in the storage and D pointer is incremented.

The manipulation of the queue is performed by two independent mechanisms, a first one being in charge of the Enqueue process and the second being in charge of the Dequeueing process. The first unit performs READ operations in the storage at addresses A(E), A(D) and WRITE operations at addresses INCR(E) and A(E) (where A(E) and A(D) stand for the address of pointers E and D in the storage). The second unit performs READ operations in the storage at addresses A(E), A(D) and INCR(D) and a write operation at address A(D).

The manner of accessing those locations is

such that no interference between the Enqueueing and Dequeueing operations is possible and each communicating entity can access the queue independently of the other entity.

Brief Description of the Drawings

FIG. 1 shows the general principle of a cyclic queue mechanism

FIG. 2 shows the different status of a cyclic queue located in memory.

FIG. 3 is a flowchart of the steps involved in processing a dequeue operation.

FIG. 4 is a flowchart of the steps involved in processing an Enqueue operation.

FIG. 5 is a flowchart of a Dequeue operation including a hysteresis function.

FIG. 6 is an overall view of a first illustrative embodiment of the invention.

FIG. 7 describes with details the first illustrative embodiment of the invention.

FIG. 8 shows examples of timing detailing the access of memory 640 through communication bus 650.

FIG. 9 is a second illustrative embodiment of the invention.

FIG. 10 shows the architecture using two queues providing a bidirectional path between two independent units.

FIG. 11 shows the architecture providing a multi-priority level communication path.

Description of the Invention

Referring to Figure 1, the general principle of a queueing mechanism is shown. Queue 120 is defined as a finite set of connected locations wrapping around from head to tail. Work Element size depends on the amount of information to be transferred in that path. The queue may be implemented by making the locations physically consecutive or by using a pointer in each location to indicate the address of the next location of the queue. The actual number of locations assigned to a queue can be set depending upon the expected maximum usage, to avoid reaching a full queue, or can be varied according to changing requirements by a monitoring function. Such locations could be mapped into 1 to N bytes of storage or 1 to N registers.

Each component, the sending component 100 and receiving component 110 manipulate queue 120 by means of Enqueue (sending component) and Dequeue (receiving component) functions.

The queue is based upon the use of a memory with a capacity for storing N work elements WE,

where N is an integer. Two pointers are defined to manage the memory.

A D pointer which points to the last WE dequeued from the queue and an E pointer which points to the last WE enqueued in the queue.

These pointers define the "active" portion of the queue - the part that contains the actual queued data. Note that these pointers must obey exactly these rules in order to enable the non-locking algorithm to work. Normal queue mechanisms usually have these pointers specifying the current Work Element to enqueue or dequeue.

FIGS 2A, 2B and 2C show the 3 possible states of queue 120.

FIG 2A shows a queue in a normal state: room is available in the queue.

FIG. 2B describes a queue having a "FULL" state. All requested "Enqueue" commands are to be refused.

FIG. 2C describes a queue having an "Empty" state.

The different states can be distinguished by the following criteria:

when E = D the queue is in an Empty state

when INCR(E) = D the queue is "full"

where INCRE(E) = E + 1(module N).

When none of the preceding relations are true the queue is in a "Normal" state, that is to say that a WE can be dequeued or a new WE can be enqueued.

The INCR operator obeys the following algorithm :

INCR is triggered with a pointer value holding the address of a queue location. When the operator completes, pointer is updated with the address of the logically following location of the queue.

In case of a linked set of locations, the linked-to pointer is returned.

In the case of a set of physically consecutive locations, the consecutive location is returned except in case of wrap around. In that case, the address of the first location in the queue is returned.

With respect to FIG. 3, a flow chart is shown describing the steps involved to process a dequeueing operation. Step 301 is performed by receiver component 110 to check whether the queue is empty or not. In case the queue is empty, the dequeue is rejected and a flag QE (waiting on queue empty) is set to warn the receiver of the status of the queue (step 302). If the queue is not empty, a work element WE is returned to the

receiver and the pointer D is updated via the INCR operator, (step 303). Then, (step 304) a QF flag (Queue Full) is tested to determine whether the queue was full before the dequeuing process. In the case QF is ON, the flag is reset (step 305). QF and QE can be used to provide a signalling function and therefore permit the sending/receiving component to suspend enqueue/dequeue when a Full/Empty condition has been detected.

Referring to FIG. 4, a flow chart is shown describing the steps involved in the process of an enqueueing operation. Step 401 is performed by the sender 100 to check whether the queue is in full state or not. In case the queue is full, the Enqueue is rejected and flag QF (Queue Full) is set to warn the sender of the status of the queue (step 402). If the queue is not full, a WE is placed in its location and the pointer is updated (step 403). Then, (step 404) the queue is tested to detect if the receiver is waiting for a transition from empty to non-empty state. In case of such a transition, flag QE is reset to inform the receiver 110 (step 405).

Optionally, in case the enqueueing of Work Elements by the sender is faster than the dequeuing by the receiver, the queue can be constantly swapping from full to non-full states. This constant swapping can be avoided by applying a Hysteresis as shown in FIG. 5 in which the solution to avoid the signalling until a predefined number (X) of dequeues have been made by the receiver. This value X could be dynamically controlled by the receiver, because this component has the knowledge of the amount of signals received from the sender. The value of X can be selected as a compromise between the amount of signalling performed and the time Work Elements may wait without being serviced as a result of the signalling delay.

Steps 501, 502 and 503 are respectively identical to steps 301, 302 and 303 of FIG. 3. In step 504, the status of the QF flag is tested. When the QF flag is OFF, the dequeuing process is completed. When the QF flag is on, the process goes on with step 505 which increments a counter. Then, (step 506) the value of that counter is tested to determine whether a predefined number (X) of sequences have been performed by the receiver since the "full" status of the queue was set. If the value of the counter equals the predetermined value X, then step 507 is performed, which resets the counter and the QF flag, and eventually, the process completes. If the value of the counter does not equal the predetermined value X then the dequeuing process is completed directly.

Figure 6 shows an overall view of an illustrative embodiment of the invention. Two communicating units, Unit(A) 100 and Unit (B) 110 share access to a storage device 640 through communication bus-

ses 650 and 660 attached to a storage controller 630. Each of the communication units A and B are attached to equivalent control blocks 610 and 620 capable of manipulating the storage queues via Enqueue and Dequeue operations. In the following, the communication path from Unit A to Unit B (A to B Queue) will be discussed but it should be noticed that the same elements can be used to provide a communication path from Unit B to Unit A or different priority level communication paths, as will be described in Figures 10 and 11.

The number of Communicating Devices can be augmented as needed without requiring any modification of the described implementation. Then, each paired Communicating Devices can implement as many paths (full-duplex if needed) between them as required by the application.

In the latter cases, the unique storage 640 houses the multiple queues.

Control blocks 610 and 620 are in charge of the following functions:

Control block 610 makes enqueues from A-B queue

Control block 620 makes dequeues from A-B queue.

Therefore, both control blocks 610 and 620 are totally independent. Control block 610 enqueues in A-B queue while control block 620 may dequeue from the same queue.

In the case where the foregoing elements are used to provide other queues, a B-A queue for instance, control blocks 610 and 620 will be in charge of the additional following functions.

Control block 610 makes dequeue from B-A queue

Control block 620 makes enqueue from B-A queue

Thus, for a given queue, the enqueueing and dequeuing processes are managed separately and independently by control block 610 and 620. The sending unit A can enqueue without interfering with the receiving B while dequeuing information.

In the following, only the mechanisms involved in the manipulation of A-B queue communication path from Unit A to Unit B will be discussed.

Referring to Figure 7, which is composed of Figure 7A and 7B, an illustrative embodiment of the present invention is described in detail.

For the sake of clarity, and to emphasize the essence of the invention, the following assumptions have been made.

It is assumed that control blocks 610 and 620 treat only one Dequeue or Enqueue command at a time, or, in other words, that each control block manipulates one queue at a time (the A to B queue

in this case). If multiple queues must be managed simultaneously, additional control blocks such as control blocks 610 and 620 can be used.

The hysteresis handling as described in Figure 5 will not be discussed in the following since the man knowledgeable in the art will easily adapt the following description to introduce the concept of hysteresis in the embodiment of the invention. The size of the WE is also assumed to be constant.

Both communication units 100 and 110 share bus 650 which includes an address bus 705, a data bus 706 and a control bus 703. Therefore, memory accesses via bus 650, through READ and WRITE operations are serialized for the purposes of this description, as will be described in Figure 8. Control block 620 is not represented but is identical to control block 610 and is connected to bus 650 and to unit B (170).

Control block 610 includes a comparator 730, the output of which is connected to a register 731. Two registers 732 and 733 are connected to data bus 706 and to the two first inputs of comparator 730. The output of register 733 is connected to a INC 734 (incremental logic), the output of which is the third input of comparator 730. The output of INC 734 also connected to the inputs of registers 735 and 736 which are respectively connected to data bus 706 and address bus 705. Control block 610 also includes a selector 737 being connected to n bi-register 760₁, 760₂...760_n. Each bi-register 760_i contains the values of A_i(E) and A_i(D) which correspond to the addresses of pointers E and D relating to the queue i (for example the indicia of A to B queue is 1, the indicia of B to A queue is 2, etc.). The value of a selected register A_i(E) or A_i(D) can be transmitted to address bus 705 via bus 738 via the selector 737. Registers 739 and 740 allow the communication of data between unit A 100 and data bus 706. All registers, selector 737, comparator 730, increment logic 734 are controlled by a control logic 741 as will be described herebelow.

Control block 620 is identical to control block 610. It includes a comparator 930, the output of which is connected to n register 931. Two registers 932 and 933 are connected to data bus 706 and to the two first inputs of comparator 930. The output of register 933 is connected to a INCR 934 also connected to the inputs of registers 935 and 936 which are respectively connected to data bus 706 and address bus 705.

Control block 620 also includes a selector 937 being connected to n bi-register 960₁...960_n. Each bi-register 960_i contains the values of A_i(E) and A_i(D) which correspond to the addresses of pointers E and D relating to the queue i. The value of n selected register A_i(E) or A_i(D) can be transmitted to address bus 705 via bus 938 and via the selec-

tor 937. Registers 939 and 940 allow the communication of data between unit B 110 and data bus 706. All registers, selector 937, comparator 930, INCR logic 934 are controlled by a control logic 941.

Enqueueing process into the A-B queue.

Before any enqueueing or dequeuing process, an initialization sequence is required (IPL, power on or reset): for each queue i inside storage device 640, the two registers 760_i (A_i(E) and A_i(D)) are filled in with the storage addresses of E and D pointers used to manipulate the i queue.

The Enqueue operation starts with a request from Unit A 100 which enables "Enqueue - CMD" 701. As soon as "C_BUS" 703, part of the Communication Bus 650, indicates that the Bus is available, Control Logic 741 raises lead "VALID A(E)" 702 which activates Selector 737. The value "QUEUE_ID" carried on bus 704 is used by this Selector to choose the proper A_i(E) register. When activated by "VALID A(E)" 702, Selector 737 puts the contents of the chosen register A_i(E) 760_i on "A_BUS" 705.

An appropriate command is put by the Control Logic 741 on "C_BUS" 703 to indicate a read operation. The Storage Device is selected and returns the read E value on "D_BUS" 706. When the "C_BUS" 703 indicates "D_BUS" 706 validity, Control Logic 741 sets lead 707 on. This lead sets Register 732. Now Register 732 holds the E value.

When the "C_BUS" 703 indicates that the Communication Bus 650 is available, Control Logic 741 raises "VALID-A(D)" 708 which activates the Selector 737. The value "QUEUE_ID" 704 is used by this Selector to elect the proper A_i(D) register. When activated by "VALID-A(D)" 708, Selector 737 puts the contents of the selected register A_i(D) 760_i on "A_BUS" 705.

An appropriate command is put by the Control Logic 741 on "C_BUS" 703 to indicate a read operation. The Storage Device is selected and returns the read D value on "D_BUS" 706. When "C_BUS" 703 indicates "D_BUS" 706 validity, Control Logic 741 sets the lead "LATCH-R1" 709 on. This lead sets the Register 733. Now Register 733 holds the D value.

The following describes the different steps to perform the Enqueue operation according to Figure 4.

Step 401: INCR(E) = D?

Since lead "ENQUEUE-CMD" 701 is on, Comparator 730 compares the output of Register 733 with the output of the Incrementation Logic 734. Lead "COMPARATOR-OUTPUT" 710 reflects the result of the comparison.

Step 402: WE is put in location INCR(E)

If the inputs of Comparator 730 (output of Register 733 and output of the Incrementation Logic 734) do not match, lead "COMPARATOR-OUTPUT" 710 is off. The Work Element, placed by the Communicating Device on the "WE_BUS" 714 must be put in the Storage Device at location INCR(E). When "C_BUS" 703 indicates that the Communication Bus is available, Control Logic 741 activates the lead "PUT_WE" 715 and the lead "ENQ-WE" 718 which enables Driver 736 and Driver 739 respectively. Therefore, the output of the Incrementation Logic 734 is gated to "A_BUS" 705 and "WE_BUS" 714 is gated to "D_BUS" 706. Control Logic 741 indicates on the "C_BUS" 703 a write operation. The storage Device is selected and stores the Work Element into storage and eventually the communication bus 650 is released.

E pointer is updated: E = INCR(E)

When "C_BUS" 703 indicates that the Communication Bus is available, Control Logic 741 starts the operation of updating the E pointer value. The leads "UPDATE_PTR" 716 and "VALID A(E)" 702 are activated. The new E value from Incrementation Logic 734 is put via Driver 735 on "D_BUS" 706 and E address is put via Selector output 738 on "A_BUS" 705. As soon as "C_BUS" 703 indicates the end of the memory update, Control Logic 741 activates the lead "PROCESS-DONE" 713 to indicate the completion of the operation to the Communicating Device. As soon as lead "ENQUEUE-CMD" 701 is dropped by the Communicating Device, the enqueue operation ends.

Step 403: Q Full is returned to the Communicating Device

If the inputs of Comparator 730 (output of Register 733 and output of the Incrementation Logic 734) match, the "COMPARATOR-OUTPUT" 710 is on. Control Logic 741 raises lead 711 which sets Register 731. The output of Register 731, lead "Q-FULL-EMPTY" 712 reflects to the Communicating

Device the Full state of that i queue. As soon as lead "ENQUEUE-CMD" 701 is dropped by Communicating Device 100, the Control Logic activates lead "PROCESS-DONE" 713 which resets Register 731 and the enqueue operation ends.

Dequeueing process of the B-A Queue.

As seen in the foregoing, the dequeuing process of the A to B queue is performed by Control Block 620. Similarly, Control Block 610 performs the dequeuing process of the B to A queue.

As pointed out above, the following operations imply as if that registers 960; are filled in with the values Ai(E) and Ai(D) during the initialization sequence.

20 E and D pointers acquisition.

When lead "DEQUEUE-CMD" 917 is raised by Communicating Device 100 to indicate a dequeue request, the dequeue operation starts. As soon as "C_BUS" 703, part of the Communication Bus 650, indicates that the Bus is available, the Control Logic 941 raises lead "VALID-A(D)" 908 which activates Selector 937. The value "QUEUE_ID" 904 is used by this Selector to choose the proper Ai(D) register. When activated by "VALID-A(D)" 908, the selector 937 puts the content of the selected register Ai(D) on the "A-BUS" 905.

An appropriate command is put by Control Logic 941 on "C_BUS" 703 to indicate a read operation. Storage Device 640 is selected and returns the read D value on "D_BUS" 706.

When the "C_BUS" 703 indicates the D_BUS 706 validity, the control logic 941 makes the lead 907 on. This lead sets the register 932. Now register 932 holds the D_value.

As soon as "C_BUS" 703 indicates that the communication bus 650 is available, the control logic 941 raises the lead "valid A(E)" 902 which activates the selector 937. The value "Queue_ID" 904 is used by this selector 937 to select the proper A_(E) register 902, the selector 937 puts the content of the selected A_(E) register on the "A_BUS" 705.

An appropriate command is put by the control logic 941 on the "C_BUS" 703 to indicate a read operation. The storage device is selected and returns the read E value on the "D_BUS" 706. Then the "C_BUS" 703 indicates the "D_BUS" 706 validity, the control logic 941 makes the lead 909 on, this lead sets the register 933. Now register

933 holds the E value.

The following describes the different steps to perform the Dequeue operation according to Figure 3.

Step 301: E = D?

Since lead "DEQUEUE-CMD" 917 is on, Comparator 930 compares the output of Register 933 with the output of the Register 932 and lead "COMPARATOR-OUTPUT" 910 reflects the result of the comparison.

Step 303: WE is read from location INCR(D)

If the inputs of Comparator 930 (Register 933 and Register 932) do not match, lead "COMPARATOR-OUTPUT" 910 is off. The Work Element must be read at address INCR(D) inside the Storage Device. When the "C_BUS" 703 indicates that the Communication Bus is available, the Control Logic 941 activates lead "PUT_WE" 915 which enables Driver 936. Then, the output of Incrementation Logic 934 is placed on the "A_BUS" 705. The control logic 941 indicates a read operation on the "C_BUS" 703. Control Logic 941 activates the lead "DEQ-WE" 919 which enables Driver 740. Therefore, the contents of "D_BUS" 706 is placed on the "WE_BUS" 914.

D pointer is updated: D = INCR(D)

When "C_BUS" 703 indicates that Communication Bus 650 is available, Control Logic 941 starts the operation of updating the D pointer value. Leads "UPDATE_PTR" 916 and "VALID-A(D)" 908 are activated. The new D value from Incrementation Logic is put via Driver 935 on "D_BUS" 706 and D address is put via Selector output on "A_BUS" 705. As soon as "C_BUS" 703 indicates the end of the memory update, Control Logic 941 activates the lead "PROCESS-DONE" 913 to indicate the completion of the operation to the Communicating Device. As soon as the lead "DEQUEUE-CMD" 917 is dropped by the Communicating Device (the dequeued Work Element has been read from "WE_BUS" 914 by the Communicating Device), the dequeue operation ends.

Step 302: Q Empty is returned to Communication Device 900

If the inputs of Comparator 930 (Registers 932 and 933) match, lead "COMPARATOR-OUTPUT"

910 is on. Control Logic 941 raises lead 911 which sets Register 931. The output of Register 931, lead "Q-FULL-EMPTY" 912, reflects to the Communicating Device the Empty state of that queue. As soon as lead "DEQUEUE-CMD" 917 is dropped by the Communicating Device, Control Logic 941 activates lead "PROCESS-DONE" 913 which resets Register 931 and the dequeue operation ends.

It should be noticed that Control Clocks 610 and 620 distinguish each other only in the fact that for a given queue i, one performs the Enqueue operation and the other is in charge of the dequeue operation. Each element of Control Block 610, referenced 7XX, corresponds exactly to its equivalent element in Control Block 620 and which is referenced 9XX.

Referring to Figure 8, examples of chronograms are shown detailing the access of storage 640 through Control Blocks 610 and 620 to perform parallel Enqueue and Dequeue operations.

Chronogram (a) relates to Control Block 610 communication with storage 640 while an Enqueueing operation (in A to B queue). Enqueueing operation includes sequences 801, 802 and 803. Sequence 801 is a "READ" (R) operation which corresponds, for instance, to the reading of the E value on "D_BUS" 706 before its storage into register 732 as described above. Similarly, sequence 802 corresponds to the reading of the D value on "D_BUS" 706 before its storage into register 733. Sequence 803 is the access to storage 640 to write the work element WE into the location INCR(E).

Chronogram (b) summarizes the access of control block 620 to storage 640 in a Dequeue operation (A to B queue). This Dequeueing operation includes sequences 811, 812 and 813. Sequence 811 is a "READ" operation which corresponds, for instance, to the reading of the D value on "D_BUS" 706 before its storage into register 932. Similarly, sequence 812 corresponds to the reading of the E value on "D_BUS" before its storage into register 933. Sequence 813 is the access to storage 640 by Control Block 620 to read the work element from location INCR(D) into storage 640.

Chronogram (c) summarizes the access to Communication Bus 650 by Control Blocks 610 and 620. The Read/Write operations are serialized at the bus/storage level but Enqueue/Dequeue operations can be overlapped (Reads and Writes are interleaved) thus taking benefit from our proposed algorithm allowing Enqueue and Dequeue simultaneously on the same queue. No interlocking (implicit or explicit) is required.

Since the foregoing implementation considers a shared Communication Bus 650, memory accesses are serialized at the READ/WRITE elemen-

tary operations. That slight disadvantage involved by the straightforward but simple embodiment of the invention is easily avoided when considering a more sophisticated architecture.

In respect to Figure 9, which is composed of Figures 9A and 9B, a second more sophisticated embodiment of the invention is shown. The parallelism for READ/WRITE Access of the shared storage 640 is fully provided by using a multiple busses structure and a multiple port memory 640. Communication A 100 access multiaccess memory 640 through Control Block 610 and a first Communication Bus 650. Similarly, Communication Unit B 110 accesses memory 640 through Control Block 620 and a second Communication Bus 660. Control Blocks 610 and 620 are in charge of the following functions:

Control Block 610 makes Dequeues from B-A queue and makes Enqueues to A-B queue.

Control Block 620 makes dequeues from A-B queue and makes Enqueues to B-A queues.

The Enqueue/Dequeue function involves the same operation as those detailed in the foregoing with respect to Figure 6. For a given queue, A-B queue for instance, Control Blocks 610 and 620 can operate independently. As a matter of fact, the Enqueue function involves steps 401, 402 and 404 (with respect of Figure 4) which performs READ operations at addresses $A_i(E)$, $A_i(D)$ and WRITE operations for addresses INCR(E) and $A_i(E)$. On the other hand, the Dequeue function involves steps 301, 302, 303 and 304 (with respect to Figure 3) which performs READ operations at addresses $A_i(E)$, $A_i(D)$ and INCR(D) and a WRITE operation at address $A_i(D)$. Therefore no location in the storage can be addressed by Control Blocks 610 and 620 both performing a WRITE operation. No interference is possible between units 100 and 110 while enqueueing and dequeuing and no interlocking is involved.

Referring to FIG. 10, a bi-directional path is shown between two independent units: a unit A 100 and a unit B 110. A queue (one queue per one way path established between the two units) is provided to permit the communication between Unit A 100 and Unit B. A queue 120 allows the transmission of information from Unit(A) 100 to Unit 110 (B) and a queue 130 allows the transmission from unit 110 (B) to unit 100 (A). Therefore a full duplex path between the two units requires two queues, one for each direction.

Referring to Fig. 11, another way of combining elementary queues to provide multiple communication paths between units 100 and 110 is shown. Queues 210 and 220 allow multiple level of priority in the communication between the two units. For

this purpose, the sender unit can specify a queue identifier indicating the desired level of priority.

Any sophisticated communication architecture can be provided by combining elementary queues. It should be noticed that all the queues can be located in the same Storage Device 640 which can provide as many communication paths (between multiple couple of paired units) as required.

Claims

1. Non-locking queueing mechanism for transferring information from a sending unit (100) to a receiving unit (110), through a queueing means (120) containing element areas in which said sending unit enqueues information elements and from which said receiving unit dequeues information elements, said element areas being located in a storage (640)

characterized in that it comprises

a first pointer register (640) for logging a first pointer (D) identifying the element area in said queueing means (120) where the last dequeued information element, if any, was located

a second pointer register (640) for logging a second pointer (E) identifying the element area in said queueing means where the last enqueued information element, if any, was located,

a first control block (610) activated by said sending unit (100) to enqueue said information element into said queueing means (120) and for updating said second pointer

a second control block (620) activated by said receiving means to dequeue said information element from said queueing means (120) and for updating said first pointer whereby said sending (100) and receiving (110) means can work independently without interlock or any serialism between said units.

2. A mechanism according to Claim 1 characterized in that said first control block (610) further comprises:

-a first incrementing means (734) receiving said second pointer register contents (E) and returning, when activated, a value identifying, in said queue means (120), the element area which follows that identified by said second pointer (E).

-a first comparator (730) for comparing the value returned by said first incremented means

- (734) with said first pointer D.
- means (735, 741) for logging the value returned by said first incrementing means into said second pointer register (E)
- said first control block (610) performs the following steps:
- comparing the value returned by said first incrementing means (734) with said first pointer (D) by said first comparator (730)
 - supplying to said sending unit (100) a signal indicating that said queueing means is full when said first comparator detects an equality when said first comparator detects an inequality:
 - placing the information element supplied by said sending unit (100) into said queueing means (120) in the element area identified by the value returned by said first incrementing means (734)
 - logging said incremented returned value into said second pointer register
- said second control block (620) further comprises
- a second comparator (930) for comparing said first and second pointer values
 - a second incrementing means (934) receiving said first pointer register contents (D) and returning, when activated, a value identifying, in said queueing means (120), the element which follows that identified by said first pointer (D)
 - means (935, 941) for logging the value returned by said second incrementing means (934) into said first pointer register (D)
- said second control block (620) performs the step of:
- comparing said first and second pointer by said second comparator (930)
 - supplying a signal to said receiving unit (110) indicating that the queueing means is empty when said second comparator (930) detects an equality when said second comparator (930) detects an inequality:
- activating said second incrementing means (934)
- 5 -providing the information element identified by the value returned by said second incrementing means to said receiving unit (110)
- 10 -logging the value returned by said second incrementing means into said first pointer register (D).
- 15 3. Mechanism according to Claim 2 characterized in that said first and second control blocks (610) and (620) associated with said storage (650) manage multiple queues between said sending unit (100) and said receiving unit (110) by providing as many pairs of first/second pointers (E,D) as there are queues.
- 20 4. Mechanism according to Claim 2 characterized in that said first and second control block (610) and (620) associated with said storage (650) manages a first (A-B) and second (B-A) queue between said two units (100) (110) providing a full duplex path between said two units (100) (110),
- 25 5. said first control block (610) makes enqueues from said first queue (A-B) and dequeues from said second queue (B-A)
- 30 said second control block (620) makes dequeues from said first queue (A-B) and enqueues from said second queue (B-A).
- 35 5. Mechanism according to Claims 3 or 4 characterized in that
- 40 6. said storage (640) is a multiple port memory having at least a first communication bus (650) and a second communication bus (660)
- 45 said first control block (610) enqueues said information element into said queueing (120) by addressing said storage (640) through said first communication bus (650)
- 50 said second control block (620) dequeues said information element into said queueing means (120) by addressing said storage (640) through said second communication bus (660)
- 55 6. Mechanism according to any one of Claims 1 to 5 characterized in that it comprises a hysteresis mechanism for avoiding too much swapping of said queueing means from full to non-full states.




FIG.1





FIG.2


FIG. 3

DEQUEUE PROCESS



ENQUEUE PROCESS

FIG.4






FIG. 5




FIG. 6

FIG. 7a




FIG. 7 b





FIG.8 A TO B QUEUE MANIPULATION



0 273 083

FIG.9a






FIG. 9b





0273 083

FIG. 10




FIG. 11 : MULTIPLE LEVELS OF PRIORITY



European Patent
Office

EUROPEAN SEARCH REPORT

Application number

EP 86 43 0058

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.4)
D, Y	WO-A-8 400 835 (WESTERN ELECTRIC CO., INC.) * Page 2, line 30 - page 3, line 35; figures 1-3 * & US-A-4 507 760 (Cat. D) ---	1-6	G 06 F 5/06
Y	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 15, no. 8, January 1973, pages 2495-2496, New York, US; BAUMER: "Read-only memory controlled buffering" * Figure; page 2495, line 1 - page 2496, line 10 *	1,2	
Y	---		
Y	DE-A-3 149 678 (SIEMENS AG) * Page 6, lines 6-26; page 9, line 33 - page 10, line 1; page 10, lines 17-26; figure 2 *	3-5	
Y	---		TECHNICAL FIELDS SEARCHED (Int. Cl.4)
Y	US-A-4 040 027 (VAN ES et al.) * Abstract *	6	G 06 F
A	---		
A	DE-A-2 714 106 (TELEFONBAU UND NORMALZEIT GmbH) * Claim; figures 1,2 *	1,2	

The present search report has been drawn up for all claims			
Place of search	Date of completion of the search	Examiner	
THE HAGUE	11-09-1987	BEINDORFF W.H.	
CATEGORY OF CITED DOCUMENTS			
X : particularly relevant if taken alone	T : theory or principle underlying the invention		
Y : particularly relevant if combined with another document of the same category	E : earlier patent document, but published on, or after the filing date		
A : technological background	D : document cited in the application		
O : non-written disclosure	L : document cited for other reasons		
P : intermediate document	& : member of the same patent family, corresponding document		